

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

APPLICATION FOR LETTERS PATENT

**Kernel Emulator
for Non-Native Program Modules**

Inventors:

Barry Bond

ATM Shafiqul Khalid

ATTORNEY'S DOCKET NO. MS1-665US

03647535-0304
FOR SE 24860

1
2 **CROSS-REFERENCE TO RELATED APPLICATION**

3 This application claims priority from U.S. Provisional Patent Application
4 Serial No. 60/244,410, titled "Kernel Emulator for Non-Native Program Modules"
5 filed on October 30, 2000.

6
7 **TECHNICAL FIELD**

8 This invention generally relates to a technology facilitating the operation of
9 non-native program modules within a native computing platform. This invention
10 further generally relates to a technology facilitating the interoperability of native
11 and non-native program modules within a native computing platform.

12
13 **BACKGROUND**

14 Applications designed to run on a specific computing platform do not
15 operate on a different computing platform. Generally, software is inextricably
16 linked to the computing platform on which it is designed to operate. Software
17 written and compiled to operate within minicomputer running a specific
18 implementation of the Unix operating system will not function within a hand-held
19 computer using a proprietary operating system.

20 A computing platform typically includes an operating system (OS) and
21 computing hardware architecture. Examples of OSs include these Microsoft®
22 operating systems: MS-DOS®, Windows® 2000, Windows NT® 4.0, Windows®
23 ME, Windows® 98, and Windows® 95. Examples of computing hardware
24 architecture include those associated with these Intel® microprocessors: 80286,
25 Pentium®, Pentium® II, Pentium® III, and Itanium™.

1 Examples of computing platforms includes 16-bit platforms (such as
2 Microsoft® MS-DOS® and Intel® 80286), 32-bit platforms (such as Microsoft®
3 Windows® NT® and Intel® Pentium® II), and 64-bit platforms (such as Intel®
4 Itanium™ and an appropriate 64-bit OS). A computing platform may also be
5 called a platform, computing environment, or environment.

6 Specific versions of applications are designed to operate under a specific
7 platform. These applications may be called “native” when they execute under their
8 specific platform. For example, Microsoft® Office 2000 is an application designed
9 to operate on 32-bit platform. In other words, Microsoft® Office® 2000 is a
10 native application relative to its 32-bit platform. However, these 32-bit
11 applications may be called “non-native” when they execute under a different
12 platform, such as a 64-bit platform.

13 An example of a program-module target platform (or simply “target
14 platform”) is the platform an executable program (e.g., program module,
15 application, program) was targeted to run. For a program module, its target
16 platform is also its native platform. For example, if one builds a Microsoft®
17 Office® application to run under Windows® 2000 32-bit X86 OS environment
18 then for that image target platform would be 32-bit x86.

19 An application program is the primary example of a “program module” as
20 the term is used herein. However, the term “program module” includes other
21 executable software that may not be labeled an application.

22 **Typical Computer Architecture**

23
24 Typical computer architecture is multi-layered. From the bottom up, it
25 includes the hardware layer, the operating system (OS) layer, and the application

layer. Alternatively, these layers may be described as the hardware layer, the kernel mode layer, and the user mode layer.

Fig 1 illustrates the layers of typical computer architecture 100. The top of the architecture is the user mode 110. It includes applications, such as applications 112a-e. These applications communicate with a set of APIs 120. Typically, this API set is considered part of the OS, and thus, part of the computing platform.

The next layer of the architecture is the kernel mode 130. This may be generally called the “kernel” of the OS. Since it is part of the OS, it is part of the computing platform.

A kernel of an OS is the privileged part of the OS—the most trusted part of the OS. It is an inner layer of code. It typically operates I/O 132, security 134, display control (i.e., access to the screen) 136, memory management 138, and other privileged functions 139. The kernel has sole access to the hardware in the hardware layer 150 via device drivers 142 and other hardware interfaces 144.

Kernel APIs 140 are those APIs within the kernel that arbitrate access to the kernel functions. The applications typically do not call the kernel directly. Instead, the applications call the APIs 120 and the APIs, in turn, may call the kernel (in particular the kernel APIs 140).

Although Fig. 1 does not show the components 132-144 of the kernel 130 with connections between them, these components are connected as is necessary. The coupling lines are omitted from the drawing for the sake of simplicity and clarity.

Below the kernel mode 130, there is the hardware layer 150. This layer includes all of the hardware of the actual computer. This includes the

processor(s), memory, disk I/O, other I/O, etc. The platform also includes the hardware layer.

Therefore, a computing platform includes the hardware layer, the kernel layer, and typically the user-mode APIs 120.

Interoperability and Compatibility

Application compatibility has been big concern since computing platforms started evolving. People want to run desired applications in their chosen platform in the ideal world. However, in the real world, it's very difficult to run an application in a different host platform that it wasn't written for. For example, 32-bit x86 application cannot run on 64-bit Merced (IA64) environment. The problem becomes worse when people buy a more powerful machine with a different platform than they used to for a long time. Immediately all the applications in the old platform becomes useless unless they find some way to use that in the new environment.

Each platform has its corresponding body of native applications that are designed to run under it. When a new generation of platform is released, software developers generally upgrade their products to run under the new generation platform. Software developers do this for many reasons, including marketing, technology, and economics.

For similar reasons, OS developers wish to make their products backwards compatible. In this way, older generations of applications may run on the latest generation of the OS (and thus the latest generation of platform). In other words, if non-native applications can run under a native platform (including the new OS), this encourages users to purchase the new OS because they are not forced to

1 discard their current applications and purchase new versions. This also gives
2 software developers time to develop upgrades to their applications.

3 Herein, an example of compatibility is a non-native program module
4 functioning appropriately and peacefully co-existing with native program modules
5 within a native computing environment (e.g., an OS).

6 As used herein, an example of interoperability is the ability of both native
7 and non-native program modules to share resources (such as access data within
8 each other's memory space or a shared memory space) and/or work together and
9 cooperatively.

10 For the sake of clarity and simplicity, an example is used herein to illustrate
11 the problem of incompatibility of non-native applications and non-interoperability
12 between native and non-native applications. The non-native program modules are
13 called 32-bit applications because they are designed to operate on a 32-bit
14 platform. The native applications are called 64-bit applications because they are
15 designed to operate on the native platform, which is 64-bit. This is provided as
16 one example and not for limitation. Those of ordinary skill in the art understand
17 and appreciate that there exists other combinations of native and non-native
18 applications and native platforms.

19 **Running Non-Native Applications on a Native Platform**

20
21 Consider this: Running non-native applications on a native platform. More
22 specifically, consider this example: Running 32-bit applications in a 64-bit
23 environment. Assume, for this example, that the 64-bit platform is an upgrade to
24 an existing popular 32-bit platform (on which the 32-bit applications are designed
25 to run).

1 One advantage of a 64-bit platform over a 32-bit platform is that the much
2 larger memory space can be addressed. 32-bit processors in a 32-bit platform can
3 address about 2GB of memory, but a 64-bit processors in a 64-bit platform can
4 address terabytes of memory.

5 One of the problems with going from 32-bit platform to 64-bit platform is
6 that the application programming interfaces (APIs) grow from 32-bit to 64-bit.
7 Therefore, the new APIs deal with larger memory pointers, larger memory
8 addressable space, etc. Thus, the 64-bit APIs are incompatible with calls from 32-
9 bit applications. Thus, software developers need to recompile their applications to
10 64-bit and release a new version for the 64-bit platform. For many large end
11 applications (particularly on servers), porting to 64-bit is the most appropriate
12 alternative to take advantage of the new capabilities of a 64-bit OS. But for most
13 applications, it is not best to port the application for numerous reasons, such as the
14 sales projections for the 64-bit platform may not enough to justify the expense of
15 porting, the application is very difficult to port, etc.

16 However, it is desirable to have the 32-bit applications operate and function
17 correctly on a 64-bit platform. In other words, it is desirable to have a non-native
18 (e.g., 32-bit) application run properly in a native (e.g., 64-bit) environment.

19 The capability of running 32-bit applications on the 64-bit platform is
20 highly desirable for easing the transition from the popular 32-bit platform to the
21 new 64-bit platform. Not only may this ease the transition, it may reduce the
22 burden on software developers to immediately port their software to the 64-bit
23 platform. It also gives the new 64-bit platform an existing library of software
24 (e.g., the old 32-bit applications). Furthermore, it is desirable to have 32-bit and
25 64-bit applications interoperate.

Virtual Machine (VM)

In similar situations in the past, the solution has nearly always been to emulate the hardware on which the non-native applications would run. In other words, the native platform emulates the non-native hardware environment (such as a 32-bit processor) in which the non-native applications (such as 32-bit applications) run. Furthermore, the non-native OS (such as a 32-bit OS) run in the non-native hardware environment because the non-native applications need the non-native OS. This non-native hardware emulation is often called a "virtual machine" (VM) and the combination of the VM and the non-native OS is sometimes called the "box" or non-native OS "box" (e.g., a MS-DOS® box).

Fig. 2 illustrates the same multi-layered architecture of Fig. 1 except it includes a virtual machine 200. The VM emulates the non-native hardware layer 250 in software. To run non-native applications (such as 212a and 212b) within a VM, a non-native OS runs on top of the emulated non-native hardware layer 250. This non-native OS includes a user mode 210 and a kernel mode 230. Non-native APIs 220 are in the user mode 210 and a non-native kernel is in the kernel mode 230. As a result, these three layers (user mode, kernel mode, and hardware) are executed solely in software in a segregated "box" 200. There are other possible arrangements for a VM. For example, some or all of the VM components may be implemented within the native kernel 130.

Notice that the actual non-native kernel 230 is being executed within the VM 200. The non-native kernel 230 is not being emulated. Rather, it is running on top of an emulated hardware layer 250 of the VM 200.

Those of ordinary skill in the art understand the VM model. Although this model is very common, it consumes a large amount of resources and processing power (i.e., it is expensive). This is because it emulates a non-native hardware and executes the complete non-native OS on top of that emulated hardware. In addition, non-native applications operating within a VM box cannot interoperate ("interop") with native applications running within the native platform. Moreover, they cannot interoperate with other non-native applications running within other VMs. Thus, VM emulation is expensive and incompatible.

SUMMARY

Described herein is a technology facilitating the operation of non-native program modules within a native computing platform. This invention further generally relates to a technology facilitating the interoperability of native and non-native program modules within a native computing platform.

Specifically, this technology involves an emulation of the kernel of the non-native operating system. Instead of interacting with the native kernel of the native computing platform, the non-native program modules interact with a non-native kernel emulator. This emulator handles the necessary conversions and translations. With this non-native kernel emulation, native and non-native program modules are interoperable. Except for the kernel emulator, none of the program module (native or non-native) and none of the other portions of the native computing platform are aware of the emulation. The computing environment and other program modules appear to be non-native to the non-native program modules. Likewise, the non-native program modules appear to be native to the computing environment and the native program modules.

1 This summary itself is not intended to limit the scope of this patent. For a
2 better understanding of the present invention, please see the following detailed
3 description and appending claims, taken in conjunction with the accompanying
4 drawings. The scope of the present invention is pointed out in the appending
5 claims.

6 **BRIEF DESCRIPTION OF THE DRAWINGS**

7
8 The same numbers are used throughout the drawings to reference like
9 elements and features.

10 Fig. 1 is a schematic block diagram showing a computing architecture.

11 Fig. 2 is a schematic block diagram showing a computing architecture with
12 an example of a virtual machine (VM) implementation.

13 Fig. 3 is a schematic block diagram showing a computing architecture
14 showing an embodiment in accordance with the invention claimed herein.

15 Fig. 4 is a schematic block diagram of a kernel emulator in accordance with
16 the invention claimed herein.

17 Fig. 5 is a flow diagram showing a methodological implementation of the
18 invention claimed herein.

19 Fig. 6 is an example of a computing operating environment capable of
20 implementing an implementation (wholly or partially) of the invention claimed
21 herein.

22 **DETAILED DESCRIPTION**

23
24 The following description sets forth specific embodiments of a kernel
25 emulator for non-native program modules here that incorporate elements recited in

1 the appended claims. These embodiments are described with specificity in order to
2 meet statutory written description, enablement, and best-mode requirements.
3 However, the description itself is not intended to limit the scope of this patent.

4 Described herein are one or more exemplary implementations of a method
5 and system of fusing portions of a print medium. The inventors intend these
6 exemplary implementations to be examples. The inventors do not intend these
7 exemplary implementations to limit the scope of the claimed present invention.
8 Rather, the inventors have contemplated that the claimed present invention might
9 also be embodied and implemented in other ways, in conjunction with other
10 present or future technologies.

11 An example of an embodiment of a kernel emulator for non-native program
12 modules here may be referred to as an “exemplary kernel emulator.”

13 **Incorporation by Reference**

14
15 This provisional application is incorporated by reference herein: U.S.
16 Provisional Patent Application Serial No. 60/244,410, titled “Kernel Emulator for
17 Non-Native Program Modules” filed on October 30, 2000.

18 **Introduction**

19
20 The one or more exemplary implementations, described herein, of the
21 exemplary kernel emulator may be implemented (whole or in part) by a kernel
22 emulator 400 and/or by a computing environment like that shown in Fig. 6.

23 The exemplary kernel emulator, described herein, provides a mechanism
24 with which non-native applications can be run on a native platform transparently.
25

For instance, 32-bit applications can be run on 64-bit platform by using the exemplary kernel emulator.

This mechanism can also be used to plug-in a CPU simulator to broaden the scope. For example, the exemplary kernel emulator enables an application to run on a computer with a Merced processor where the application was written for x86 instruction set.

Using an exemplary implementation of the kernel emulator, the kernel of nearly any platform may be emulated and different CPU instructions can be simulated by plug-in CPU simulators.

Herein, references to "kernel emulation" mean emulation of a non-native kernel within the context of a native platform. Similarly, references, herein, to "kernel emulator" mean a non-native kernel emulator within the context of a native platform.

Overview of Kernel Emulation

As described in the above Background section, traditional solution for enabling non-native programs to operate on top of a native platform is virtual machine (VM). However, kernel emulation is less expensive than VMs. Less expensive in terms resources dedicated to emulation (e.g., resources include processor cycles, memory, overhead, etc.). Moreover, unlike VM emulation, all applications (including both native and non-native) may interoperate ("interop").

In this exemplary kernel emulation, described herein, the non-native applications believe that they are running on an operating system (OS) with their non-native kernel. Instead, their non-native kernel is being emulated. The non-

native kernel emulator intercepts kernel calls made by the non-native applications and translates them into native kernel calls.

The exemplary kernel emulator is an emulator of the kernel of a non-native OS. It allows non-native applications to run within a native environment by emulating the non-native kernel of a non-native OS. The hardware is not emulated. Nor are the APIs. Instead, the kernel of the non-native OS is emulated. The emulated kernel translates non-native calls (from the applications and their APIs) into calls that can be interpreted and handled by the native OS. Moreover, the emulated kernel interprets and handles information flowing to the non-native applications from the native OS.

To accomplish this, an exemplary kernel emulator may perform one or more of the following functions:

- translate hardware instruction set from non-native to native
- intercept kernel calls to convert from non-native memory addressing to native memory addressing
- constrain memory access of non-native applications to only the memory space addressable by such applications (e.g., the lower 2GB of RAM). By doing so, the native memory manager need not be bothered with such constraints
- manage memory shared by both non-native and native applications
- accommodate a suitable CPU simulator that can be plugged in to simulate a variety of CPU architectures.

Computer Architecture Implementing Exemplary Kernel Emulator

Fig. 3 illustrates a computer architecture 300 within which the exemplary kernel emulator may be implemented. The architecture includes a native platform, which includes hardware 350, kernel mode 330, and set of native APIs 320 of user mode 310. The top of the architecture 300 is the user mode 310. It includes applications. More specifically, it includes native applications (such as applications 312a-c) and non-native applications (such as applications 314a and 314b).

The native applications communicate with a set of native APIs 320. Typically, this native API set 320 is considered part of the native OS, and thus, part of the native computing platform. The non-native applications communicate with a set of non-native APIs 322. Typically, this non-native API set 320 is considered part of a non-native OS, and thus, the set is not part of the native computing platform.

Without the exemplary kernel emulation, the non-native applications (such as 314a and 314b) would not function. Why? They cannot call the native APIs 320 because they are unaware of them and because the native APIs are written to understand the instruction set and memory organization of the native platform, which are different from the instruction set and memory organization of the non-native platform expected by the non-native applications. In addition, if the non-native applications called their own non-native APIs 322, the non-native APIs would fail for the same reasons given above. Furthermore, the non-native APIs attempts to call the kernel APIs would call native kernel APIs 340. Such calls would not be understood by the native kernel APIs 340.

The non-native APIs 322 include “stub” code that directs their calls to a non-native kernel, which does not exist in the native platform. However, this “stub” code is the bridge between the “user mode” and the “kernel mode.” In this exemplary kernel emulator, this stub code is modified so that it calls a non-native kernel emulator 400.

The next layer of the architecture 300 is the native kernel mode 330. This may be generally called the native “kernel” of the native OS. Since it is part of the native OS, it is part of the native computing platform.

This native kernel 330 includes the native kernel APIs 340 and various kernel functions 333, such as I/O, security, display control (i.e., access to the screen), memory management, and other privileged functions. The native kernel 330 has sole access to the hardware in the hardware layer 350 via device drivers 342 and other hardware interfaces 344. Native kernels APIs 340 are those APIs within the kernel that arbitrate access to the kernel functions.

Although Fig. 3 does not show the components 333-344 of the kernel 330 with connections between them, these components are connected as is necessary. The coupling lines are omitted from the drawing for the sake of simplicity and clarity.

Below the kernel mode 330, there is the hardware layer 350. This layer includes all of the hardware of the actual computer. This includes the processor(s), memory, disk I/O, other I/O, etc. The native platform also includes the hardware layer.

Therefore, native computing platform includes the hardware layer 350, the kernel layer 330, and typically the user-mode native APIs 320. The native

1 platform does not include the non-native APIs 322 or the non-native kernel
2 emulator 400.

3 4 **Exemplary Kernel Emulator**

5 As shown in Fig. 3, the non-native kernel emulator 400 is part of both the
6 user layer of the user mode 310 and the kernel layer of the native kernel mode
7 330. Consequently, portions of the non-native kernel emulator 400 may operate
8 primarily in the user layer and others may operate primarily in the kernel layer.
9 Alternatively, the non-native kernel emulator 400 may be part of just the kernel
10 layer. The non-native kernel emulator 400 communicates with non-native APIs
11 332, the native kernel, and directly with the memory of the hardware layer 350.

12 Although Fig. 3 shows the non-native kernel emulator 400 within the kernel
13 mode layer, it may be implemented in the user mode layer in alternative
14 embodiments. It may also be implemented partially in the user mode and partially
15 in the kernel mode.

16 Fig. 4 illustrates the non-native kernel emulator 400, which is an
17 implementation of the exemplary kernel emulator. It includes an emulator 410 of
18 non-native kernel APIs, which may also be called the non-native CPU simulator
19 (or simply the “CPU simulator”). To avoid confusion with nomenclature, this will
20 be called CPU simulator 410.

21 Although this name seems to imply hardware emulation, the hardware is
22 not being emulated by the CPU simulator 410 or by any other portion of the non-
23 native kernel emulator 400. As discussed in the Background section, the OS is
24 inextricably linked to the hardware (included the CPU). Therefore, the native and
25 non-native hardware is a major consideration in non-native kernel emulation.

1 Those of ordinary skill in the art see and understand that this CPU simulation is
2 fundamentally different from an emulation of hardware (like in the VM model).

3 The exemplary kernel emulator may have CPU simulators for different
4 platforms. The exemplary kernel emulator finds a CPU simulator for the target
5 CPU architecture (of the non-native application) and uses for its kernel emulation.

6 The CPU simulator 410 includes a translator 412 for translating non-native
7 to native CPU instructions sets, a converter 414 for converting non-native to
8 native word sizes, and a memory constrainer 416 to limit the memory accessible to
9 non-native applications to that which such applications are capable of addressing.

10 In the exemplary implementation, the translator 412 and converter 414
11 operate primarily within the user mode. The advantage to running them in user
12 mode is that bugs in that code will crash only the user-mode process. If code
13 running in kernel mode crashes, it will likely crash the entire machine.

14 Calls to the Kernel by Non-Native Applications

15
16 The CPU simulator 410 receives the kernel calls from non-native
17 applications. More specifically, the CPU simulator 410 receives the kernel calls
18 from non-native APIs.

19 The kernel emulator 400 does not perform the functions of a kernel.
20 Rather, it translates the non-native calls from their non-native format to a native
21 format. It then passes the translated calls onto the native kernel for the native
22 kernel to process as if the call was from a native application. However, to do this
23 may involve more than a simply translation of a language. It may involve a
24 conversion of paradigms between the non-native and native platforms.
25

For example, within a 32-bit platform (such as Microsoft® Windows® 98 and an x86 instruction-set processor), when an application wants to make a kernel API call, it pushes its arguments onto a stack (within memory). The arguments sit in memory. If the application wants to call an API and pass it “3” and a pointer to a string of text, then the application writes “3” to the memory (onto the stack) and then writes the address of the string to memory, as well. The API can read these values from memory to pick up the arguments when the API is executed.

Continuing the example, on a 64-bit platform (such as a 64-bit OS and an Intel® Itanium™ processor), argument values are passed within the CPU itself. More specifically, the argument values are passed within the registers of the CPU. Instead of an application storing the argument values in a stack (i.e., memory), it stores them in CPU registers.

The translator 412 performs instruction emulation. It handles the emulation of the non-native processor’s instruction set.

The converter 414 manages both argument format and word size. It changes the argument format from non-native to native and changes the word size from non-native to native.

Together, the translator 412 and converter 414 manage the argument-passing-via-stack convention of the 32-bit platform to the argument-passing-via-registers convention of the 64-bit platform. This is an example of a paradigm change in how argument values are passed. The CPU simulator 410 accounts for this change with its translation of the CPU instruction set by translator 412.

In another example, in a 32-bit platform, the word size is 32 bits long. Thus, addresses and data are typically written in 32-bit sized words in memory. In

a 64-bit platform, the word size is 64 bits long. Thus, addresses and data are typically written in 64-bit sized words in memory.

Parameters (i.e., arguments) are typically one word long. Therefore, the parameters for the 64-bit kernel APIs are 64-bits long; rather than 32-bits long for the 32-bit kernel APIs.

The converter stretches the 32-bit arguments passed by the non-native applications (and APIs) from 32-bits to 64-bits. One example of how it does this is by padding the argument with 32 leading zeroes. In the exemplary implementation, the converter 414 operates primarily within the user layer.

The exemplary kernel emulator, described herein, assumes that the word size for the non-native platform is smaller than the addressable memory space for the native platform. In an alternative embodiment, this assumption may be reversed. Those of ordinary skill in the art understand that this would a reversal of the techniques described herein.

As discussed above, differing platforms typically have differing word sizes. Consequently, the size of the addressable memory space differs, as well. Therefore, the converter 414 converts the addresses (in particular, pointers) as part of its word conversion.

For instance, the memory addresses of a 32-bit platform may be up to 32-bits long; thus, the maximum addressable memory is approximately 4GB (if one of the bits is reserved, then the maximum is about 2GB). The memory addresses of a 64-bit platform may be up to 64-bits long; thus, the maximum addressable memory is measured in terabytes. Thus, the size of addressable memory differs between the 32-bit and 64-bit platforms.

Constraining Memory for Non-Native Applications

The native OS is unaware that non-native applications are non-native. Each non-native application appears to be a native process because the kernel emulator 400 presents it to the native OS (specifically, the kernel) as a native application.

Native applications can access the full range of addressable memory. However, non-native applications cannot. Instead, they are limited to a subset of memory available to a native application. For example, a 32-bit application is forced to operate within the lower 2GB of the terabytes of potential memory space allocated to a native application.

Since the native memory manager is unaware of a distinction between native and non-native applications, the memory constrainer 416 of the non-native kernel emulator 400 forces this limitation.

When a non-native application requests a memory workspace, this request goes through the kernel emulator 400 rather than the native kernel. The memory constrainer 416 of the kernel emulator 400 reserves all of the available memory except for the lower portion. This unreserved portion matches the memory space addressable by the non-native application. This unreserved portion is available for the non-native application operating within.

For example, if non-native applications can only address a maximum of 2GB and if 8TB (terabytes) are allocated to typical native application, then the memory constrainer 416 will reserves all of the memory from 2GB to 8TB. That reserved memory is unavailable to the non-native applications. When the native memory manager allocates memory for the non-native applications, it will allocate memory from the lower 2GB range.

Continuing with this example, when a non-native application wants memory, it calls the kernel emulator 400, and the emulator calls the native memory manager. The native memory manager allocates memory from the lower 2GB because that is all that is available. The native memory manager is unaware that the application seeking the memory is non-native. The manager need not know because the emulator 400 assures that the manager can only allocate memory within the proper range for the non-native application.

Intercepting Kernel Calls

Applications do not typically call the kernel directly. Rather, the applications call the APIs, which in turn, call the kernel.

With the exemplary kernel emulator, the non-native APIs call the kernel emulator 400 rather than calling the actual native kernel. The kernel emulator understands the calls from the non-native APIs and it is designed to handle them. As described above, it converts the non-native API calls into native calls to the actual native kernel.

Instruction Set Translation

The native and non-native platforms typically have differing CPU instruction sets. For example, the instruction set for an x86-based hardware is different from the instruction set for the Intel® Itanium™ hardware. A different instruction set implies a different machine code; a different assembly language; and often there are different operational paradigms.

1 The translator 412 of the kernel emulator 400 of Fig. 4 translates non-native
2 CPU instruction sets to native CPU instruction sets. In the exemplary
3 implementation, the translator 412 operates primarily within the user layer.

4 When comparing the instructions of any one CPU instruction set to that of
5 another CPU instruction set, many of the instructions are functionally similar if
6 not identical. It may take multiple instructions in one set to accomplish the same
7 task as one instruction in another task. Nevertheless, most instructions in one set
8 can be translated into instructions in another set.

9 10 Communication with the Native Kernel

11 Fig. 4 shows that the non-native kernel emulator 400 also includes an I/O
12 unit 430 for communicating with the native kernel. Via the I/O unit, converted
13 kernel calls are sent to the native kernel, memory allocation requests are made to
14 the native kernel, and all other communication to the native kernel is performed.
15 In addition, communication from the native kernel is received by the I/O unit.

16 17 Shared Memory

18 As shown in Fig. 4, the non-native kernel emulator 400 includes a shared
19 memory facilitator 450 coordinate memory sharing between native and non-native
20 applications. In the exemplary implementation, the shared memory facilitator 450
21 operates primarily within the user layer.

22 Although they run in separate allocated memory spaces, some native
23 applications actually share memory between processes. In other words,
24 applications running in their own isolated memory space share some of the
25 memory (even when they have different addresses) with other applications so that

the same memory is visible within multiple processes. For example, when an application writes a value into the memory within its own address space, it instantly appears in the memory space of a separate application. And vice versa. This is a high performance technique for applications to share data since it need only be written once, but it appears to and is instantly available to multiple applications.

An example of data that is shared across multiple applications in a 32-bit OS (like Microsoft® Windows® 98) includes a list of open windows (titles and position).

The key problem with the shared memory for non-native applications is that the content of the shared memory often includes pointers. The native applications use native memory addressing (e.g., 64-bit long addresses) and thus, read and write in the native addressing in the shared memory. However, the non-native applications don't understand the native addressing because they expect their non-native format (e.g., 32-bit long addresses).

One approach to solving this problem is to redesign some of the APIs so that it is aware that it is using shared memory. Thus, the APIs know how to read and write memory using native addressing format.

Another approach, and the one implemented by the shared memory facilitator 450, is to make copies of the shared memory. Thus, the original is in a native addressing format and the copy is in non-native addressing format. The two versions are regularly synchronized. This is done for the TEBs and PEBs structures.

An example of a "process" includes the abstraction of environment to execute instruction in the target application to get some job done. Process control

block or process environment block (PEB) may have necessary information to describe the process in the native platform and in the target environment. A PEB may be generically called a process shared data structure (SDS).

The PEB is a fixed-formatted structure of memory that each process has containing environmental information about the process, such as such as command line for the process, info about what files it has open, its position on the display, and the like.

The format of the PEB for native applications is different from the format for non-native applications. In the kernel emulation, there are two PEBs per application—a native PEB and a non-native PEB. This includes both native and non-native applications. Whenever a non-native application attempts to access the PEB, it accesses the non-native version. Whenever the native kernel attempts to access the PEB, it accesses the native version. These two versions are regularly synchronized. The address conversion (between native and non-native addressing formats) is done during the synchronization.

An example of a “thread” includes a path of execution with in a process. Thread control block or thread environment block (TEB) may have the necessary information to describe a thread in the host and in the target environment. A TEB may be generically called a thread shared data structure (SDS).

The TEB is similar to the PEB. It is a fixed-formatted structure of memory that stores information about a thread running inside of a process, such as last error value from an API call, pointer to the stack for the thread, and the like. Each thread has a TEB structure associated with it.

Similarly to the PEB, the format of the TEB for native applications is different from the format for non-native applications. In the kernel emulation,

1 there are two TEBs per thread—a native TEB and a non-native TEB. This
2 includes both native and non-native applications. Whenever a non-native
3 application attempts to access the TEB, it accesses the non-native version.
4 Whenever the native kernel attempts to access the TEB, it accesses the native
5 version. These two versions are regularly synchronized. The address conversion
6 (between native and non-native addressing formats) is done during the
7 synchronization.

8 9 **Methodological Implementation of the Exemplary Kernel Emulator**

10 Fig. 5A shows methodological implementation of the exemplary kernel
11 emulator performed by the kernel emulator 400 (or some portion thereof). This
12 methodological implementation may be performed in software, hardware, or a
13 combination thereof.

14 At 510 of Fig. 5, an application is loaded (i.e., initialized). At 512, if the
15 application is determined to be a native application, then is processed as normal, at
16 530, by the OS. At 512, if the application is determined to be a non-native
17 application, then, at 514, it is determined the target (i.e., non-native) platform and
18 the CPU simulator for that platform is selected and implemented. Alternatively,
19 steps 512 and 514 are unnecessary when the target platform may be fixed.

20 At 516 of Fig. 5A, the exemplary kernel emulator constrains the memory to
21 that which is addressable by the non-native application. At 518, the exemplary
22 kernel emulator establishes data structures, which are copies of the PEBs and
23 TEBs, but they are in a non-native format. At 519, the exemplary kernel emulator
24 asks the CPU simulator to start executing instructions in host program. At 520,
25 this method of initiating a non-native application ends. The non-native application

continues to run. Fig. 5B illustrates the steps performed by the kernel emulation while the non-native application is running.

Fig. 5B shows methodological implementation of the exemplary kernel emulator performed by the kernel emulator 400 (or some portion thereof). This methodological implementation may be performed in software, hardware, or a combination thereof.

At 550 of Fig. 5B, the exemplary kernel emulator translates the calls to a non-native kernel into calls to a native kernel. It does so, for example, in the manner described herein. It translates platform paradigms and CPU instruction sets. It converts word sizes—including converting addresses.

At 560, the exemplary kernel emulator synchronizes the native and non-native TEBs and PEBs. The exemplary kernel emulator continues to do these steps as long as the non-native application executes.

Additional Details

In general, while attempting to run an application, a loader creates a native process using that application and transfers control of execution to some entry point in the process. It creates the PEB and the TEB for the initial thread. The loader is the part of the OS that initializes an application for it to run.

While creating the native process, the loader checks the type of application. If the type of application is not native, it instead creates the kernel emulator process (kernel-emu-process) and transfers the control to the entry point in kernel-emu-process. The kernel-emu-process is designed in such a way that non-native applications feel that it's running in the native environment it was written for. Therefore, exemplary kernel emulator provides a synthesized non-native platform

for the non-native application. Whenever non-native applications need to communicate with the OS, kernel-emu-process just proxies that request. The native platform is happy because it thinks its running a native process, because kernel-emu-process don't violate any constraint may required native processes.

At the time of loading a non-native application, exemplary kernel emulator creates the necessary environment for the application so that other programs running in the host machine can identify this process as if a native application is running. The exemplary kernel emulator also creates the necessary environment for the target platform inside the host (i.e., native platform) so that the running application feels that it's running in the native platform.

When creating a kernel-emu-process, the shared memory facilitator 450 generates a target (i.e., non-native) Process Environment Block (target PEB) and native kernel generates a native PEB. There is some mapping between the native PEB and the target PEB. In the same way, the shared memory facilitator 450 generates a target (i.e., non-native) Thread Environment Block (target TEB) and native kernel generates a native TEB appropriate for the running a thread in the process. There is some mapping between the native TEB and the target TEB.

Other native processes may get information from the native PEB and TEB while querying for a kernel-emu-process and thread. The native PEB and TEB in a kernel-emu-process work as a proxy.

If there are any other necessary environment settings, exemplary kernel emulator may create a version for the target and keep some mappings between the two.

Exemplary Computing System and Environment

Fig. 6 illustrates an example of a suitable computing environment 900 within which an exemplary kernel emulator, as described herein, may be implemented (either fully or partially). The computing environment 900 may be utilized in the computer and network architectures described herein.

The exemplary computing environment 900 is only one example of a computing environment and is not intended to suggest any limitation as to the scope of use or functionality of the computer and network architectures. Neither should the computing environment 900 be interpreted as having any dependency or requirement relating to any one or combination of components illustrated in the exemplary computing environment 900.

The exemplary kernel emulator may be implemented with numerous other general purpose or special purpose computing system environments or configurations. Examples of well known computing systems, environments, and/or configurations that may be suitable for use include, but are not limited to, personal computers, server computers, thin clients, thick clients, hand-held or laptop devices, multiprocessor systems, microprocessor-based systems, set top boxes, programmable consumer electronics, network PCs, minicomputers, mainframe computers, distributed computing environments that include any of the above systems or devices, and the like.

Exemplary kernel emulator may be described in the general context of computer-executable instructions, such as program modules, being executed by a computer. Generally, program modules include routines, programs, objects, components, data structures, etc. that perform particular tasks or implement

particular abstract data types. Exemplary kernel emulator may also be practiced in distributed computing environments where tasks are performed by remote processing devices that are linked through a communications network. In a distributed computing environment, program modules may be located in both local and remote computer storage media including memory storage devices.

The computing environment 900 includes a general-purpose computing device in the form of a computer 902. The components of computer 902 can include, by are not limited to, one or more processors or processing units 904, a system memory 906, and a system bus 908 that couples various system components including the processor 904 to the system memory 906.

The system bus 908 represents one or more of any of several types of bus structures, including a memory bus or memory controller, a peripheral bus, an accelerated graphics port, and a processor or local bus using any of a variety of bus architectures. By way of example, such architectures can include an Industry Standard Architecture (ISA) bus, a Micro Channel Architecture (MCA) bus, an Enhanced ISA (EISA) bus, a Video Electronics Standards Association (VESA) local bus, and a Peripheral Component Interconnects (PCI) bus also known as a Mezzanine bus.

Computer 902 typically includes a variety of computer readable media. Such media can be any available media that is accessible by computer 902 and includes both volatile and non-volatile media, removable and non-removable media.

The system memory 906 includes computer readable media in the form of volatile memory, such as random access memory (RAM) 910, and/or non-volatile memory, such as read only memory (ROM) 912. A basic input/output system

1 (BIOS) 914, containing the basic routines that help to transfer information
2 between elements within computer 902, such as during start-up, is stored in ROM
3 912. RAM 910 typically contains data and/or program modules that are
4 immediately accessible to and/or presently operated on by the processing unit 904.

5 Computer 902 may also include other removable/non-removable,
6 volatile/non-volatile computer storage media. By way of example, Fig. 6
7 illustrates a hard disk drive 916 for reading from and writing to a non-removable,
8 non-volatile magnetic media (not shown), a magnetic disk drive 918 for reading
9 from and writing to a removable, non-volatile magnetic disk 920 (e.g., a "floppy
10 disk"), and an optical disk drive 922 for reading from and/or writing to a
11 removable, non-volatile optical disk 924 such as a CD-ROM, DVD-ROM, or other
12 optical media. The hard disk drive 916, magnetic disk drive 918, and optical disk
13 drive 922 are each connected to the system bus 908 by one or more data media
14 interfaces 926. Alternatively, the hard disk drive 916, magnetic disk drive 918,
15 and optical disk drive 922 can be connected to the system bus 908 by one or more
16 interfaces (not shown).

17 The disk drives and their associated computer-readable media provide non-
18 volatile storage of computer readable instructions, data structures, program
19 modules, and other data for computer 902. Although the example illustrates a hard
20 disk 916, a removable magnetic disk 920, and a removable optical disk 924, it is to
21 be appreciated that other types of computer readable media which can store data
22 that is accessible by a computer, such as magnetic cassettes or other magnetic
23 storage devices, flash memory cards, CD-ROM, digital versatile disks (DVD) or
24 other optical storage, random access memories (RAM), read only memories
25 (ROM), electrically erasable programmable read-only memory (EEPROM), and

1 the like, can also be utilized to implement the exemplary computing system and
2 environment.

3 Any number of program modules can be stored on the hard disk 916,
4 magnetic disk 920, optical disk 924, ROM 912, and/or RAM 910, including by
5 way of example, an operating system 926, one or more application programs 928,
6 other program modules 930, and program data 932. Each of such operating
7 system 926, one or more application programs 928, other program modules 930,
8 and program data 932 (or some combination thereof) may include an embodiment
9 of an interceptor, a call-converter, an argument converter, a translator, a shared-
10 memory manager, an instruction-translator, an address-translator, a memory
11 constrainer, a target-platform determiner, an instruction-type detector, a translator
12 selector, a target-platform simulator.

13 A user can enter commands and information into computer 902 via input
14 devices such as a keyboard 934 and a pointing device 936 (e.g., a "mouse").
15 Other input devices 938 (not shown specifically) may include a microphone,
16 joystick, game pad, satellite dish, serial port, scanner, and/or the like. These and
17 other input devices are connected to the processing unit 904 via input/output
18 interfaces 940 that are coupled to the system bus 908, but may be connected by
19 other interface and bus structures, such as a parallel port, game port, or a universal
20 serial bus (USB).

21 A monitor 942 or other type of display device can also be connected to the
22 system bus 908 via an interface, such as a video adapter 944. In addition to the
23 monitor 942, other output peripheral devices can include components such as
24 speakers (not shown) and a printer 946 which can be connected to computer 902
25 via the input/output interfaces 940.

Logical connections between computer 902 and the remote computer 948 are depicted as a local area network (LAN) 950 and a general wide area network (WAN) 952. Such networking environments are commonplace in offices, enterprise-wide computer networks, intranets, and the Internet.

When implemented in a LAN networking environment, the computer 902 is connected to a local network 950 via a network interface or adapter 954. When implemented in a WAN networking environment, the computer 902 typically includes a modem 956 or other means for establishing communications over the wide network 952. The modem 956, which can be internal or external to computer 902, can be connected to the system bus 908 via the input/output interfaces 940 or other appropriate mechanisms. It is to be appreciated that the illustrated network connections are exemplary and that other means of establishing communication link(s) between the computers 902 and 948 can be employed.

In a networked environment, such as that illustrated with computing environment 900, program modules depicted relative to the computer 902, or portions thereof, may be stored in a remote memory storage device. By way of example, remote application programs 958 reside on a memory device of remote computer 948. For purposes of illustration, application programs and other

executable program components such as the operating system are illustrated herein as discrete blocks, although it is recognized that such programs and components reside at various times in different storage components of the computing device 902, and are executed by the data processor(s) of the computer.

Computer-Executable Instructions

An implementation of an exemplary kernel emulator may be described in the general context of computer-executable instructions, such as program modules, executed by one or more computers or other devices. Generally, program modules include routines, programs, objects, components, data structures, etc. that perform particular tasks or implement particular abstract data types. Typically, the functionality of the program modules may be combined or distributed as desired in various embodiments.

Exemplary Operating Environment

Fig. 6 illustrates an example of a suitable operating environment 900 in which an exemplary kernel emulator may be implemented. Specifically, the exemplary kernel emulator(s) described herein may be implemented (wholly or in part) by any program modules 928-930 and/or operating system 926 in Fig. 6 or a portion thereof.

The operating environment is only an example of a suitable operating environment and is not intended to suggest any limitation as to the scope or use of functionality of the exemplary kernel emulator(s) described herein. Other well known computing systems, environments, and/or configurations that are suitable for use include, but are not limited to, personal computers (PCs), server

1 computers, hand-held or laptop devices, multiprocessor systems, microprocessor-
2 based systems, programmable consumer electronics, wireless phones and
3 equipments, general- and special-purpose appliances, application-specific
4 integrated circuits (ASICs), network PCs, minicomputers, mainframe computers,
5 distributed computing environments that include any of the above systems or
6 devices, and the like.

7 **Computer Readable Media**

8
9 An implementation of an exemplary kernel emulator may be stored on or
10 transmitted across some form of computer readable media. Computer readable
11 media can be any available media that can be accessed by a computer. By way of
12 example, and not limitation, computer readable media may comprise “computer
13 storage media” and “communications media.”

14 “Computer storage media” include volatile and non-volatile, removable and
15 non-removable media implemented in any method or technology for storage of
16 information such as computer readable instructions, data structures, program
17 modules, or other data. Computer storage media includes, but is not limited to,
18 RAM, ROM, EEPROM, flash memory or other memory technology, CD-ROM,
19 digital versatile disks (DVD) or other optical storage, magnetic cassettes, magnetic
20 tape, magnetic disk storage or other magnetic storage devices, or any other
21 medium which can be used to store the desired information and which can be
22 accessed by a computer.

23 “Communication media” typically embodies computer readable
24 instructions, data structures, program modules, or other data in a modulated data
25

1 signal, such as carrier wave or other transport mechanism. Communication media
2 also includes any information delivery media.

3 The term “modulated data signal” means a signal that has one or more of its
4 characteristics set or changed in such a manner as to encode information in the
5 signal. By way of example, and not limitation, communication media includes
6 wired media such as a wired network or direct-wired connection, and wireless
7 media such as acoustic, RF, infrared, and other wireless media. Combinations of
8 any of the above are also included within the scope of computer readable media.

9 **Conclusion**

10 Although the invention has been described in language specific to structural
11 features and/or methodological steps, it is to be understood that the invention
12 defined in the appended claims is not necessarily limited to the specific features or
13 steps described. Rather, the specific features and steps are disclosed as preferred
14 forms of implementing the claimed invention.
15
16
17
18
19
20
21
22
23
24
25